

CUSTOMIZING THE SPEAKING STYLE OF A SPEECH SYNTHESIZER BASED ON SEMANTIC ANALYSIS

Background of the Invention

[0001] The present invention relates generally to text-to-speech synthesis, and more particularly, to a method for customizing the speaking style of a speech synthesizer based on semantic analysis of the input text.

[0002] Text-to-speech synthesizer systems convert character-based text into synthesized audible speech. Text-to-speech synthesizer systems are used in a variety of commercial applications and consumer products, including telephone and voicemail prompting systems, vehicular navigation systems, automated radio broadcast systems, and the like.

[0003] Prosody refers to the rhythmic and intonational aspects of a spoken language. When a human speaker utters a phrase or sentence, the speaker will usually, and quite naturally, place accents on certain words or phrases, to emphasize what is meant by the utterance. In contrast, text-to-speech synthesizer systems can have great difficulty simulating the natural flow and inflection of the human-spoken phrase or sentence. Consequently, text-to-speech synthesizer systems incorporate prosodic analysis into the process of rendering synthesizer speech. Although prosodic analysis typically involves syntax assessments of the input text at a very granular level (e.g., at a word or sentence level), it does not involve a semantic assessment of the input text.

[0004] Therefore, it is desirable to provide a method for customizing the speaking style of a speech synthesizer based on semantic analysis of the input text.

Summary of the Invention

[0005] In accordance with the present invention, a method is provided for customizing the speaking style of a speech synthesizer. The method includes: receiving input text; determining semantic information for the input text; determining a speaking style for rendering the input text based on the semantic information; and customizing the audible speech output of the speech synthesizer based on the selected speaking style.

[0006] For a more complete understanding of the invention, its objects and advantages, refer to the following specification and to the accompanying drawings.

Brief Description of the Drawings

[0007] Figure 1 is a flowchart illustrating a method for customizing the speaking style of a speech synthesizer based on long-term semantic analysis of the input text in accordance with the present invention;

[0008] Figure 2 is a block diagram depicting an exemplary text-to-speech synthesizer system in accordance with the present invention; and

[0009] Figure 3 is block diagram depicting how global prosodic settings are applied to phoneme data by an exemplary prosodic analyzer in accordance with the present invention.

Detailed Description of the Preferred Embodiments

[0010] Figure 1 illustrates a method for customizing the speaking style of a speech synthesizer based on semantic analysis of the input text. While the following description is provided with reference to customizing the speaking style of the speech synthesizer, it is readily understood that the broader aspects of the present invention includes customizing other aspects of the text-to-speech synthesizer system. For instance, the expression of a talking head (e.g., a happy talking head) or the screen display of a multimedia user interface may also be altered based on the semantic analysis of the input text.

[0011] First, input text is received at step 12 into the text-to-speech synthesizer system. The input text is subsequently analyzed to determine semantic information at step 14. Semantic analysis of the input text is preferably in the form of topic detection. However, for purposes of the present invention, semantic analysis refers to various techniques that may be applied to input text having three or more sentences.

[0012] Topic detection may be accomplished using a variety of well known techniques. In one preferred technique, topic detection is based on the frequency of keyword occurrences in the text. The topic is selected from a list of anticipated topics, where each anticipated topic is characterized by a list of

keywords. To do so, each keyword occurrence is counted. A topic for the input text is determined by the frequency of keyword occurrences and a measure of similarity between the computed keyword occurrences and the list of pre-selected topics. An alternative technique for topic detection is disclosed in U.S. Patent No. 6,104,989 which is incorporated by reference herein. It is to be understood that other well known techniques for topic detection are also within the scope of the present invention.

[0013] A speaking style can impart an overall tone and better understanding of a communication. For instance, if the topic is news, then the speaking style of a news anchorperson may be used to render the input text. Alternatively, if the topic is sports, then the speaking style of a sportscaster may be used to render the input text. Thus, the selected topic is used at step 16 to determine a speaking style for rendering the input text. In a preferred embodiment, the speaking style is selected from a group of pre-determined speaking styles, where each speaking style is associated with one or more of the anticipated topics.

[0014] It is envisioned that semantic analysis may be performed on one or more subsets of the input text. For example, large blocks of input text may be further partitioned into one or more context spaces. Although each context space preferably includes at least three phrases or sentences, semantic analysis may also occur at a more granular level. Semantic analysis is then performed on each context space. In this example, a speaking style may be selected for each context space.

[0015] Lastly, the audible speech output of the speech synthesizer is customized at step 18 based on the selected speaking style. For instance, a news anchorperson typically employs a very deliberate speaking style that may be characterized by a slower speaking rate. In contrast, a sportscaster reporting the exciting conclusion of a sporting event may employ a faster speaking rate. Different speaking styles may be characterized by different prosodic attributes. As will be more fully described below, the prosodic attributes for a selected speaking style are then used to render audible speech.

[0016] An exemplary text-to-speech synthesizer is shown in Figure 2. The text-to-speech synthesizer 20 is comprised of a text analyzer 22, a phonetic analyzer 24, a prosodic analyzer 26 and a speech synthesizer 28. In accordance with the present invention, the text-to-speech synthesizer 20 further includes a speaking style selector 30.

[0017] In operation, the text analyzer 22 is receptive of target input text. The text analyzer 22 generally conditions the input text for subsequent speech synthesis. In a simplistic form, the text analyzer 22 performs text normalization which involves converting non-orthographic items in the text, such as numbers and symbols, into a text form suitable for subsequent phonetic conversion. A more sophisticated text analyzer 22 may perform document structure detection, linguistic analysis, and other known conditioning operation.

[0018] The phonetic analyzer 24 is then adapted to receive the input text from the text analyzer 22. The phonetic analyzer 24 converts the input text into corresponding phoneme transcription data. It is to be understood that

various well known phonetic techniques for converting the input text are within the scope of the present invention.

[0019] Next, the prosodic analyzer 26 is adapted to receive the phoneme transcription data from the phonetic analyzer 24. The prosodic analyzer 26 provides a prosodic representation of the phoneme data. Similarly, it is to be understood that various well known prosodic techniques are within the scope of the present invention.

[0020] Lastly, the speech synthesizer 28 is adapted to receive the prosodic representation of the phoneme data from the prosodic analyzer 26. The speech synthesizer renders audible speech using the prosodic representation of the phoneme data.

[0021] To customize the speaking style of the speech synthesizer 28, the text analyzer 22 is further operable to determine semantic information for the input text. In one preferred embodiment, a topic for the input text is selected from a list of anticipated topics as described above. Although determining the topic of the input text is presently preferred, it is envisioned that other types of semantic information may be determined for the input text. For instance, it may be determined that the input text embodies dialogue between two or more persons. In this instance, different voices may be used to render the text associated with different speakers.

[0022] A speaking style selector 30 is adapted to receive the semantic information from the text analyzer 22. The speaking style selector 30 in turn determines a speaking style for rendering the input text based on the semantic

information. In order to render the input text in accordance with a particular speaking style, each speaking style is characterized by one or more global prosodic settings (also referred to herein as "attributes"). For instance, a happy speaking style correlates to an increase in pitch and pitch range with an increase in speech rate. Conversely, a sad speaking style correlates to a lower than normal pitch realized in a narrow range and delivered at a slow rate and tempo. Each prosodic setting may be expressed as a rule which is associated with one or more applicable speaking styles. One skilled in the art will readily recognize other types of global prosodic settings may also be used to characterize a speaking style. The selected speaking style and associated global prosodic settings are then passed along to the prosodic analyzer 26.

[0023] Global prosodic settings are then applied to phoneme data by the prosodic analyzer 26 as shown in Figure 3. In a preferred embodiment, the global prosodic settings are specifically translated into particular values for one or more of the local prosodic parameters, such as pitch, pauses, duration and volume. The local prosodic parameters are in turn used to construct and/or modify an enhanced prosodic representation of the phoneme transcriptions data which is input to the speech synthesizer. For instance, an exemplary global prosodic setting may be an increased speaking rate. In this instance, the increased speaking rate may translate into a 2ms reduction in duration for each phoneme that is rendered by the speech synthesizer. The speech synthesizer then renders audible speech using the prosodic representation of the phoneme

data as is well known in the art. An exemplary speech synthesizer is disclosed in U.S. Patent No 6,144,939 which is incorporated by reference herein.

[0024] The foregoing discloses and describes merely exemplary embodiments of the present invention. One skilled in the art will readily recognize from such discussion, and from accompanying drawings and claims, that various changes, modifications, and variations can be made therein without departing from the spirit and scope of the present invention.